

Development Of A Massively Parallel NOGAPS Forecast Model

Tom Rosmond
Naval Research Laboratory
Monterey, CA 93943-5502
phone: (831) 656-4736 fax: (831) 656-4769
e-mail: rosmond@nrlmry.navy.mil
Award #: N0001498WX30167
<http://www.nrlmry.navy.mil>
mmmRos1.doc

LONG-TERM GOALS

Develop an advanced global atmospheric forecast system designed to exploit massively parallel processor (MPP), distributed memory computer architectures. Future increases in computer power from MPP's will allow substantial increases in model resolution, more realistic physical processes, and more sophisticated data assimilation methods, all of which will improve operational numerical weather predictions and provide better simulations of the Earth's climate.

OBJECTIVES

The current Navy operational global atmospheric prediction system (NOGAPS 4.0) is a highly optimized Fortran code designed to run on parallel vector, shared memory machines (CRAY's). The immediate objective of the project is to redesign the model's numerical algorithms and data structures to allow efficient execution on MPP architectures and clusters of shared memory processors. Message passing (MPI) is the paradigm chosen for communication between distributed memory processors. This work is supported by ONR Marine Meteorology, PE 0602435N (035-71).

APPROACH

Use integrations of the current operational NOGAPS as control runs to ensure reproducibility of results with the newly designed Fortran 90 code. Design efficient spectral transform algorithms for both shared memory and distributed memory architectures. For distributed memory architectures use message passing library modules in communication intensive spectral transforms and horizontal interpolation routines.

The current NOGAPS spectral formulation requires global communication for the spherical harmonic transforms. An attractive alternative is the use of quasi-uniform icosahedral grids based on local basis functions that are less communication intensive. A development effort on this next-generation NOGAPS has begun.

WORK COMPLETED

The complete NOGAPS spectral forecast model has been ported to a scalable architecture design using MPI as the communication methodology. The code has been run successfully on the Cray T3E, SGI Origin 2000, DEC 8400 SMP, and Cray C90. The computational core of the model, including the

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 1998		2. REPORT TYPE		3. DATES COVERED 00-00-1998 to 00-00-1998	
4. TITLE AND SUBTITLE Development Of a Massively Parallel NOGAPS Forecast Model				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Research Laboratory,Monterey,CA,93943-5502				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES See also ADM002252.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 4	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

solution of the dynamical equations and the diabatic processes, scales very well to at least 200 processors. The I/O and communication intensive pre-processing and post-processing parts of the model do not scale well, and are candidates for separation from the main model and ported to shared memory systems more suited for this kind of computation. Such a heterogeneous computing environment is inevitable for large, complex production codes such as NOGAPS.

A shallow water version of the icosahedral grid NOGAPS has been completed and extensively tested.

RESULTS

The scalable NOGAPS MPI code has been run extensively on the T3E at a variety of resolutions and over a varying number of processors to test performance and robustness. The figure below shows results for a T159L32 NOGAPS, which is representative of the current operational resolution, and processor numbers from 15 to 240. The computational core of the model, represented by 'diabat', 'dry_dyn', and 'MPI_trans', scale reasonably well over this range. The 'dry_dyn' line is from the dynamical equations and shows the effects of varying cache reuse, but overall is almost parallel with the perfect scaling line, i.e., the one marked 'perfect'. The 'diabat' line shows excellent scaling to 60 processors, but falls off above. This is almost entirely due to severe load imbalances in the convective parameterization, which is concentrated in the tropics.

The lines 'hist_writ', 'dagnos', and 'p2sig', are I/O, diagnostics, and pre-processing, respectively. The show total lack of scaling. For the first two, this is not surprising, since they are largely single processor operations that potentially can be improved with the incorporation of the yet to be released MPI-2 standard. The pre-processing step, however, is clearly a problem area. The cubic spline interpolations used are communication intensive and scale disastrously above 60 processors. NOGAPS post-processing shows similar properties. Current NOGAPS pre- and post-processing is clearly not appropriate for distributed memory architectures and will require alternate strategies in future heterogeneous computing environments.

The icosahedral grid NOGAPS has been used to evaluate the relative merits of local finite element and local spectral element methods on these kind of quasi-regular grids. A number of papers and presentations on the results have been published.

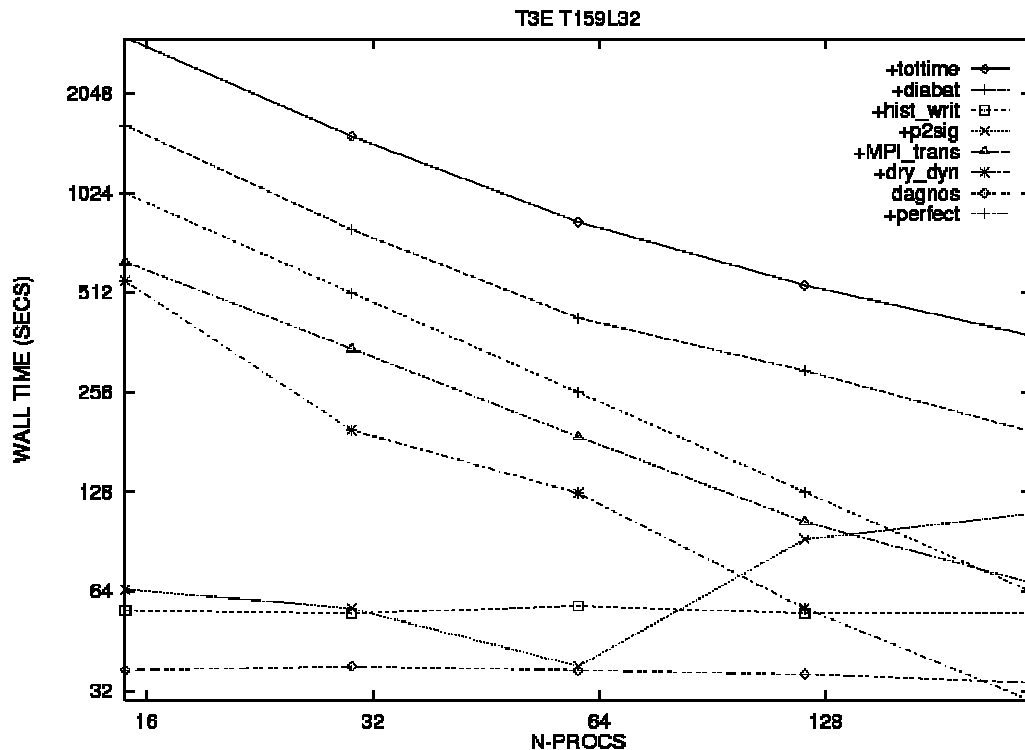
IMPACT

NOGAPS is run operationally by FNMOC and is the heart of the Navy's operational weather prediction support to nearly all DOD users worldwide. It is also run by many NRL and other Navy researchers to study atmospheric dynamics, and atmosphere/ocean interaction. Our work here targets the next generation of this system for the next generation of computer architectures. These architectures are expected to be distributed memory, commodity based systems with enormous theoretical computational power. However, exploiting this capability will require drastically redesigning many important model algorithms.

TRANSITIONS

Improved algorithms for model processes will be transitioned to 6.4 (PE 0603207N) as they are ready, and will ultimately be transitioned to FNMOC with future NOGAPS upgrades. Development of the MPI NOGAPS code has necessitated close examination of the algorithms used in the operational model, and in some cases uncovered design weaknesses and bugs that are being promptly corrected in the operational NOGAPS.

The scalable NOGAPS MPI code has been provided to FNMOC as the flagship benchmark code for their planned FY99 procurement of a scalable system to replace the current operational C90's.



RELATED PROJECTS

(1) NOGAPS 4.0 Evaluation (X0513-01): Advanced development and transition of the NOGAPS 4.0 forecast model to operational status at FNMOC. (2) The DOD CHSSI Scaled Software algorithm development for meteorological models (HPCM-96-032): Development of numerical algorithms appropriate for massively parallel computer architectures. These algorithms will be critical for inter-processor communication dependent and computationally intensive model processes.

REFERENCES

- Giraldo, F. X., 1998: Efficiency and Accuracy of Lagrange-Galerkin Methods on Unstructured Adaptive Grids. *Mathematical Modelling & Scientific Computing*, Vol. 8.
- Rosmond, T. E., 1998: A Scalable Version of the Navy Operational Global Atmospheric Prediction System Spectral Forecast Model. Submitted to *Scientific Programming*. Oct. 1998.
- Rosmond, T. E., 1998: Investigation of the Scalability of Pre-processing, Post processing, and Computational Modules of the NOGAPS Spectral Forecast Model. Presented at the ECMWF workshop: "Toward Teracomputing – The use of parallel processors in Meteorology". 16-20 November, 1998.

IN-HOUSE/OUT-OF-HOUSE RATIOS

100% in-house, 0% out-of-house.